

An Automated Occultation Network for Gravitational Mapping of the Trans-Neptunian Solar System

Daniel C. H. Gomes  and Gary M. Bernstein 

Department of Physics & Astronomy, University of Pennsylvania, 209 S. 33rd St., Philadelphia, PA 19104, USA; dchgomes@sas.upenn.edu

Received 2024 October 21; revised 2024 November 25; accepted 2024 December 14; published 2025 January 22

Abstract

We explore the potential of an array of $\mathcal{O}(100)$ small fixed telescopes, aligned along a meridian and automated to measure millions of occultations of Gaia stars by minor planets, to constrain gravitational signatures from a ‘‘Planet X’’ mass in the outer solar system. The accuracy of center-of-mass tracking of occulters is limited by photon noise, uncertainties in asteroid shapes, and Gaia’s astrometry of the occulted stars. Using both parametric calculations and survey simulations, we assess the total information obtainable from occultation measurements of main-belt asteroids (MBAs), Jovian Trojans, and trans-Neptunian objects (TNOs). We find that MBAs are the optimal target population due to their higher occultation rates and abundance of objects above Legacy Survey of Space and Time detection thresholds. A 10 yr survey of occultations by MBAs and Trojans using an array of 200 40 cm telescopes at 5 km separation would achieve 5σ sensitivity to the gravitational tidal field of a $5M_{\oplus}$ Planet X at 800 au for $>90\%$ of potential sky locations. This configuration corresponds to an initial cost of $\approx \$15$ million. While the survey’s sensitivity to tidal forces improves rapidly with increasing number of telescopes, sensitivity to a Planet X becomes limited by degeneracy with the uncertain masses of large moonless TNOs. The 200-telescope survey would additionally detect ≈ 1800 TNO occultations, providing detailed shape, size, and albedo information. It would also measure the Yarkovsky effect on many individual MBAs, measure masses of many asteroids involved in mutual gravitational deflections, and enable better searches for primordial black holes and departures from general relativity.

Unified Astronomy Thesaurus concepts: Planetary dynamics (2173); Stellar occultation (2135); Astrometry (80); Asteroids (72); Trans-Neptunian objects (1705); Gravitation (661)

1. Introduction

Stellar occultations, i.e., transits of an occulter object that temporarily blocks the light from a background star, provide a convenient method for studying physical properties and orbital dynamics of solar system objects. Historically, potential uses of the method have long been discussed, as seen in Leonhard Euler’s application for determining longitude with stellar occultations by the Moon (L. Euler 1749) and Rev. John Michell’s mention of measuring stellar angular diameters using occultations by Venus (J. Michell 1767). The 20th century witnessed an increased interest in the study of planetary atmospheres via occultations, building upon theoretical development by A. Pannekoek (1903) and C. Fabry (1929). Toward the end of the century, successful observations allowed the discovery of Uranus’ rings (J. L. Elliot et al. 1977) and the detection of Pluto’s atmosphere (W. B. Hubbard et al. 1988).

The advent of high-precision stellar catalogs has expanded possibilities in the field, especially in the study of small solar system bodies. These measurements allow determination of asteroid shapes, which can potentially constrain solar system formation models and provide better orbital tracking than attainable through direct observation of the target. Recent effort in understanding the outer regions of the solar system has motivated the prediction and observation of trans-Neptunian

object (TNO) occultations through projects such as the RECON network (M. W. Buie & J. M. Keller 2016).

Targeted occultation surveys rely on previous knowledge of each target’s orbit with precision comparable, at least, to the size of the telescope array. J. L. Ortiz et al. (2019) note that, though Gaia astrometry is at microarcsecond level, many TNO orbits are still uncertain to 300 mas. At a distance of 40 au, this corresponds to ~ 9000 km shadow path uncertainties, which implies a low probability of successful event prediction. Some objects, however, have been better constrained: the New Horizons extended mission target Arrokoth was successfully observed in three occultation campaigns (M. W. Buie et al. 2020) following orbital determination from Hubble Space Telescope astrometry, and other TNOs closely observed by the New Horizons spacecraft also had their orbital accuracy improved (S. B. Porter et al. 2022). As the Legacy Survey of Space and Time (LSST) commences and proceeds over its ~ 2025 –2035 period, nearly all of the minor bodies it detects will accrue sufficiently accurate orbital data to predict occultation tracks to $\mathcal{O}(100)$ km, hugely increasing the efficiency of targeted occultation measurements.

Main-belt asteroids (MBAs) and Jovian Trojans are more accessible targets, since they are closer to us and have larger apparent paths in the sky. These populations are also of interest in outer solar system studies, particularly due to their potential use as test particles in the search for distant gravitational signals, which could provide constraints on different ‘‘unseen planet’’ hypotheses (generally referred to as Planet X), such as the Planet Nine hypothesis evoked to explain a perceived orbital clustering of TNOs (C. A. Trujillo & S. S. Sheppard 2014; K. Batygin & M. E. Brown 2016; M. E. Brown & K. Batygin 2016;



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

S. S. Sheppard & C. Trujillo (2016). M. Rice & G. Laughlin (2019) suggested a network of ~ 2000 telescopes spread over the continental United States measuring occultations of Jovian Trojans, with the goal of detecting perturbations by a putative Planet X. Inspired by this idea, in a previous study (D. C. H. Gomes et al. 2023) we explored the ability of current and planned LSST astrometric data to constrain the presence of a distant tidal field, finding that LSST astrometry of Jovian Trojans does not significantly improve Planet X mass uncertainties. The possible contribution of an occultation array, however, remained unexplored beyond Rice and Laughlin's initial foray.

In the current paper, we propose a fixed array of telescopes aligned roughly along a meridian, automated to detect targeted occultations of Gaia stars by asteroids from any of the three main populations—MBAs, Jovian Trojans and TNOs. Starting with LSST-level astrometry for initial orbit determination, the array would narrow down center-of-mass (CM) positions through multiple occultation events per asteroid over a period of ≈ 10 yr. We explore which array design and target choices would yield better constraints on the tidal fields of a distant point mass.

This paper will focus on optimizing a telescope array for minor-body astrometry and tidal-field detection, but an array of $\gtrsim 100$ agile telescopes could also advance many other science goals. The astrometric data would also improve the solar system ephemeris generally, constrain the flux of asteroid-mass primordial black holes (T. X. Tran et al. 2024), and vastly improve knowledge of the sizes of asteroids and their nongravitational forces that drive much migration. A review of knowledge gained on TNOs through occultations is given by B. Sicardy et al. (2024), and in Section 5 we show that such an array would measure the shapes and albedos of $\approx 100 \times$ more TNOs than have been characterized to date. Individual stations would be available for other observations when they are not active for an occultation, enabling high-duty-cycle monitoring for asteroid impacts, supernova shock breakouts, and many other transient or variable phenomena occurring on timescales or cadences inaccessible to LSST.

2. Hardware Characteristics and Costs

We will assume that observations are made by an array of $N \gtrsim 100$ telescopes, each with diameter $d \sim 0.5$ m, separated by distance $D \sim 1$ km, spanning a length of $L = ND \gtrsim 100$ km (symbols and relevant default values are listed in Table 1). The array should be oriented roughly north–south, perpendicular to the typical (ecliptic) apparent motion of the occulters, but it is not necessary to have precisely spaced or oriented stations in the array. We will assume throughout that occultations are measured through a filter centered at $\lambda_0 = 600$ nm, spanning $400 \text{ nm} < \lambda < 800 \text{ nm}$, which is a common choice for maximizing signal-to-noise ratios (S/Ns) in the presence of atmospheric and zodiacal backgrounds. We will further assume a nominal survey duration of $T = 10$ yr.

Other important characteristics of the telescopes include the overall quantum efficiency (QE) η of the atmosphere/telescope/aperture/detector combination, for which we assume 50%. The latitude, maximum observable air mass X_{\max} , and duty cycle f_{duty} are important to determining the yield of the occultations, where the latter is the fraction of occultations that pass over (or under) the array that can actually be observed. We will default these to $\pm 30^\circ$, 2.3, and 0.21, respectively. For an observatory at latitude $\pm 30^\circ$, f_{duty} is limited to < 0.32 by the fraction of near-ecliptic sky that is visible at air mass $< X_{\max} = 2.3$. We reduce this by a factor

0.65 to allow for losses due to clouds, downtime, and bright twilight or moonlight that might be present—our simulations described in Section 4 will explicitly check each potential event against solar, lunar, and weather conditions rather than assuming an f_{duty} value.

The detectors on each telescope will need to read out the flux of the occulted star at $\gtrsim 200$ Hz in order to resolve the ≈ 10 ms transition from full to zero shadow in egress/ingress. Read noise below $1e^-$ will be needed. A potential solution is a single-channel photon-counting device, such as an avalanche photodiode. This detector would be paired with a CCD or CMOS detector with a field of view of several arcminutes, sufficient to find several stars in the Gaia catalogs for any pointing in the sky. This “finder” detector would be read out, the image analyzed to determine the current telescope pointing, and then the pointing adjusted to place the target star onto the photon counter's aperture. Thus, the pointing system needs to be capable of subarcsecond accuracy only for differential movements of a few arcminutes. The telescope mount requires only several-arcminute accuracy for longer slews.

The occultation events will last only $\lesssim 1$ s for MBAs and Trojans, and up to ≈ 1 minute for TNOs. We thus envision a very rapid automated cadence, in which each station can independently slew to a target, expose and analyze a registration image, shift the target onto the photon counter, and record several seconds of data, then move on to the next target in a total of $\lesssim 2$ minutes. It is possible that subarray readout of a CMOS “finder” detector could have low enough noise to obviate the need for a photon counter.

To minimize costs, each station should be as autonomous, reliable, and inexpensive as possible. We envision each station being powered by a solar array and battery, and transmitting at least station-keeping information by cellular links, so no utility hookups are needed. If cellular transmission of the observing data is too expensive, these data could be stored on local physical media that are collected periodically by maintenance crews. A modular telescope/mount/enclosure unit should be small enough that several can be placed on a truck, so that a crew traveling along the array can swap out units and bring them back to a central lab for repairs and maintenance; and similarly for computing and power systems.

2.1. Cost Equation

We will make very crude estimates of array capital costs as

$$\text{cost} = N \left[C_s + C_t \left(\frac{d}{0.5 \text{ m}} \right)^2 \right], \quad (1)$$

where C_s is a fixed cost per station for site preparation, enclosure, detectors, power, computing, and communications, and C_t is a telescope cost, which we assume will scale with the collecting area. Some rough guesses are as follows:

1. *Telescope.* Ritchey–Chretien telescopes from Planewave Instruments have costs that scale roughly linearly with collecting area for $d \lesssim 0.7$ m, with a telescope+mount cost of $\approx \$60,000$ for $d = 0.5$ m.¹ We thus adopt a default of $C_t = \$60,000$. Ordering at >100 -unit scale might make this and other components substantially cheaper.

¹ <https://planewave.com/collections/all-telescopes/>

Catadioptric telescopes are less expensive and available for $d \lesssim 0.4$ m.

2. *Detectors.* A CMOS or CCD “finder” with $\approx 10'$ FOV would not require particularly large format or high performance, so commercial off-the-shelf hardware would suffice. Commercial single-channel photon-counting devices should work well too; we want to push for high QE, but the timing requirements (millisecond levels) and count rates (100–1000 Hz) that we require are more relaxed than typical specifications. The two detectors might cost $\approx \$15,000$ per station total.
3. *Enclosure.* This would need to be designed and replicated; a simple hinged-top box would be sufficient, and would not need to be large enough to contain a human—the sides could open for maintenance. Perhaps this could be done for $\$10,000$ per station.
4. *Computing, solar power, and cellular data link.* These would cost a few thousand dollars per station.
5. *Site preparation.* This would comprise a concrete pad and a security fence and cameras, and cost perhaps $\$10,000$.
6. Least predictable would be the costs of negotiating and acquiring the leases for the sites. We ignore this important element for now, because any believable estimate would depend heavily on the choice of site and require research well outside of our astronomical expertise.

In sum, a crude guess is hence $C_s = \$40,000$ and $C_t = \$60,000$, leading to a $\$10,000,000$ cost for a 100-station array. We have neglected design costs (including software), which probably are independent of the array specifications. The above costs are all for capital outlays. There would of course be ongoing operating costs for maintenance/personnel. At most observatories these are 5%–10% of capital cost annually, accumulating to be similar to capital costs over the project lifetime. However, this unusual observatory might be more expensive because of its geographically distributed nature, or less expensive because its replication of components makes maintenance more efficient. An engineering and cost study is needed.

3. Parametric Analyses

In this section, we describe a simplified parametric estimation of the cost and performance of an occultation array, which we then use to study the constraints and optimization on array parameters to produce useful information on gravitational perturbations subject to cost bounds. In Section 5, we will reevaluate the optimization and performance using more thorough simulations of survey scenarios, guided by the results of these simpler parametric analyses.

The forecast of constraints on a given gravitational signature requires answers to the following questions:

1. What is the uncertainty in position σ_x of the target minor body attained from a single occultation event?
2. How many events (at what uncertainties) would the array be able to attain on a typical object?
3. How many test particles (occulters) can be tracked from the selected small-body population?
4. What constraints does the suite of observations place upon the gravitational anomaly?

For a given observing scenario and target population, a useful figure of merit is the total positional information σ_{tot} , defined by

$$\sigma_{\text{tot}}^{-2} = \sum_{i \in \text{targets}} \sum_{j \in \text{occultations of } i} \frac{1}{\sigma_{x,ij}^2}, \quad (2)$$

with $\sigma_{x,ij}^2$ being the measurement uncertainty on the position of a test particle i during its j th occultation event. The resulting σ_{tot} can be compared to the few-meter scale of information from ranging to spacecraft orbiting Jupiter and Saturn, or few-centimeter scale of information on the range to Mars, keeping in mind that the displacements caused by an anomalous tidal force on a test body will scale roughly as $\Delta x \propto a_T^4$, where a_T is the target's semimajor axis.

3.1. Test Particle Populations

We examine as target populations each of the three most populous reservoirs in the solar system: the MBAs, the Jovian Trojan asteroids, and the TNOs. In this section, we consider members of each population to share a common semimajor axis $a_T = 2.6, 5.2$, and 42 au for the MBAs, Trojans, and TNOs, respectively. Aside from a_T , the other critical characteristic of the target population is $N_T (> d_T)$, the number of members larger than diameter d_T . Figure 1 plots the values we have assumed for each population. Simple power laws are good fits to the regime where the MBAs and Trojans cataloged by the MPC are largely complete.² For fainter MBAs, we adopt the formula used in the initial LSST predictions (LSST Science Collaboration et al. 2009), and for Trojans we extrapolate the power-law fit. For TNOs, we adopt the rolling power law for dN/dH found by P. H. Bernardinelli et al. (2025) to fit the $6 < H_r < 8.2$ range. We forecast that the known populations in future years will be dominated by objects detected by LSST. The expected 5σ detection levels for single-epoch LSST exposures in nominal conditions (dark, zenith, median seeing) are equivalent to $m_V = 24.4$ mag, which we hence consider a bound on the usable occulter population.³ There are >10 times more usable MBAs than Trojans, and potentially 10^2 more MBAs than usable TNOs—but the physical displacements induced on a target by a tidal force scale as a_T^4 (for fixed number of orbits), so it is not a priori obvious which population will be most informative.

3.2. Measurement Errors

There are three important sources of uncertainty in the transverse position of the CM of a given target during an occultation, which we will call *shape noise* (σ_{SN}), *photon noise* (σ_γ), and *Gaia* uncertainties σ_G .

The ultimate limit to the determination of orbits through occultations is the shape noise, our ignorance of the target's projected CM location relative to the measured chords. Even perfect knowledge of the silhouette of the target would yield an uncertain CM estimate because of ignorance of the depth dimension, and variations in internal density cause additional CM shifts relative to the geometric centroid. In Section 4.2.1, we justify an estimate that $\sigma_{\text{SN}} = \alpha_S d_T$, with a nominal

² <https://www.minorplanetcenter.net/data>

³ <https://pstn-054.lsst.io/>

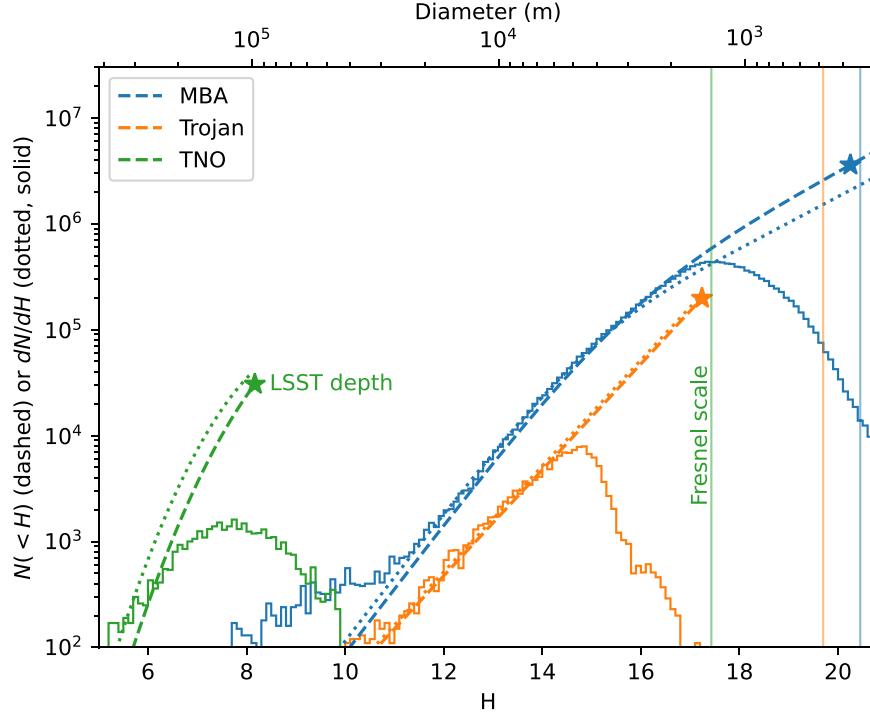


Figure 1. The dashed and dotted lines show the assumed cumulative and differential counts, respectively, for the MBA, Trojan, and TNO populations. The histograms show the differential counts dN/dm of objects currently listed by the Minor Planet Center. Each cumulative curve terminates at a star marking an apparent magnitude $m_V = 24.4$, which is roughly the limit at which objects can be detected in single visits by LSST. The vertical lines mark the Fresnel scale for each population. Note that our models are not intended to match the excesses of larger bodies.

$\alpha_S = 0.05$ giving the rms ratio of the along-track CM distance from the center of a single chord to the mean radius of the target. The cross-track uncertainty from a single-chord CM is larger, with $\sigma \approx \min(d_T, D)/\sqrt{12}$, since a simple estimator is that the CM has a uniform probability of being somewhere in a region of width $\min(d_T, D)$ centered on the observed chord. In this section, we will examine only the along-track information; the full simulation will use cross-track and multichord information, as well.

For σ_G , we use the estimated accuracy of Gaia Data Release 5 (DR5) astrometric solutions provided by the ESA.⁴

Appendix B derives the uncertainty σ_γ in the center of a chord due to Poisson noise in the arrival of photons from both the occulted star (having unocculted photodetection rate n_*) and some foreground/background count rate n_b due to atmospheric and zodiacal foreground signals, as well as sunlight reflected from the occulting asteroid. The background ratio is defined as $B = n_b/n_*$. For $B \rightarrow 0$, bright stars, and large occulters, this uncertainty becomes $\sigma_\gamma = F/\sqrt{2n_*t_F}$, where $F = \sqrt{a_T \lambda_0/2}$ is the *Fresnel scale*, which characterizes the distance over which the occulter's shadow transitions from fully bright to fully dark. This transition region takes a time $t_F = F/v_\perp$ to sweep over the telescope, where v_\perp is the relative speed of the occulter and the observer transverse to the line of sight. The Fresnel scales F and t_F are ≈ 340 m and 11 ms, respectively, for typical MBAs, growing as $\sqrt{a_T}$ for more distant occulter samples.

This best-case result is degraded if:

1. The occulter has d_T below the Fresnel scale, which leads to shallow occultations that are harder to measure, tending toward $\sigma_\gamma \propto d_T^{-5/2}$. This essentially limits the available gravitational information to objects with $d_T > F = 340(a/2.6 \text{ au})^{1/2}$ m.
2. The star is too faint to deliver at least several photons during t_F , or the background is $B \gtrsim 0.1$ and similarly degrades the S/N during the ingress/egress transitions.
3. The occulter becomes large and bright enough that its reflected light fills in the occultation light curve significantly.

Equation (B4) gives an approximation to σ_γ that spans the different regimes of source brightness, background level, and occulter size. An important consideration is that the occultation must be long and deep enough to be detected in the first place. We will require that $d_T > 4\sigma_\gamma$ to approximate this condition.

Figure 2 plots the resultant expected uncertainties σ_γ , σ_G , and σ_{SN} versus source Gaia magnitude m_G for MBAs of a few different sizes, to illustrate the domains in which each is important. One conclusion is that telescopes should have $d_T \lesssim 0.5$ m; larger telescopes would lower σ_γ , but then σ_G would dominate the total σ_x quadrature sum, negating the gain of the larger, more expensive telescopes.

3.3. Event Rates

An estimate of the number of observed occultations of stars with $m_G < m_{G,\text{max}}$ per target asteroid during the survey is obtained by finding the sky area subtended by the outline of the asteroid's track across the sky, multiplying by the sky density of potential stellar backlights and by the duty cycle of occultation observations:

⁴ <https://www.cosmos.esa.int/web/gaia/science-performance>

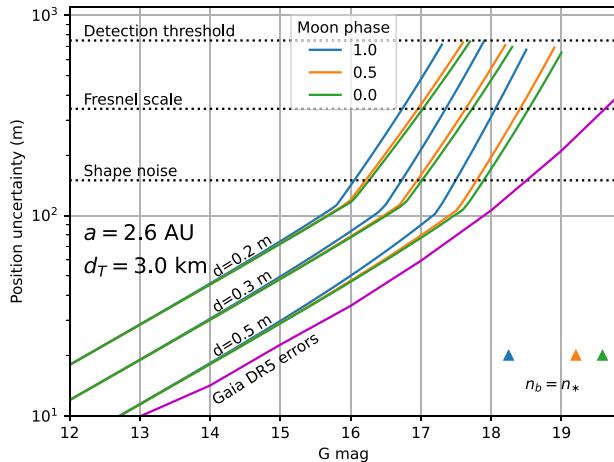
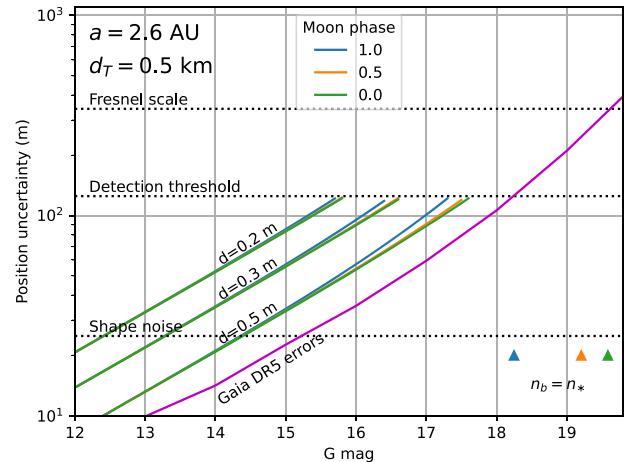


Figure 2. Sources of uncertainty in the along-track asteroid position from a single occultation chord plotted vs. magnitude of the occulted star. At left is for an occulter diameter $d_T = 3$ km, at right $d_T = 0.5$ km, both sources at a distance of 2.6 au. Each set of lines shows values of the photon noise σ_γ for new, quarter, and full Moons, and three different telescope diameters d are plotted in each case. The magenta curve plots the expected final Gaia positional uncertainties σ_G , and the horizontal line marked “shape noise” marks the floor σ_{SN} set by the unknown shape and density distribution of the occulter. The “detection threshold” marks the uncertainty $\sigma_\gamma = d_T/4$ that we consider an upper limit for detectable occultations, and the σ_γ curves terminate at this level. The Fresnel scale is also plotted, and the small triangles mark the G magnitudes at which background counts equal star counts for different moon phases. For the 0.5 km occulter, near the diffraction scale, almost all occultations are limited by photon noise, and the sky brightness is unimportant. For the larger occulter, fainter Gaia stars yield detectable occultations, and the shape noise limits the measurement accuracy for much of the detectable regime.



$$\langle N_e \rangle = f_{\text{duty}} \times n_G(< m_{G,\text{max}}) \times T \langle \text{apparent motion} \rangle \left(\frac{\text{effective array length}}{a_T} \right)$$

$$\approx 10 \left(\frac{T}{10 \text{ yr}} \right) \left(\frac{f_{\text{duty}}}{0.21} \right) 10^{0.3(m_{G,\text{max}} - 18)} \times \left(\frac{\sin^{-1}(1 \text{ au} / a_T) / a_T}{\sin^{-1}(1 / 2.6) / 2.6} \right) \left(\frac{N \times \min(d_T, D)}{200 \text{ km}} \right).$$

The final expression makes use of the following estimates. First, we estimate the mean sky density of useful Gaia stars within $\pm 10^\circ$ of the ecliptic as $n_G(< m_G) \approx 3500 \times 10^{0.3(m_G - 18)} \text{ deg}^{-2}$. This estimate does not count as “useful” those Gaia stars in crowded low-Galactic-latitude regions.

Second, we approximate the average annual length of a target’s apparent sky track, constrained to those portions with solar elongation $\beta > \beta_{\min} = 50^\circ$, as $2 \sin(1 \text{ au} / a_T)$, which is accurate to $< 10\%$ for circular orbits.

Not all positional information on a target can be used to constrain unexpected gravitational accelerations. There are 6 degrees of freedom in each target’s initial state vector, which must also be solved using the occultation data. Previous data (nominally from LSST) would be available too, but will probably be $\gtrsim 10 \times$ less accurate than the occultation data and hence not useful at occultation-survey accuracy—otherwise we would not be doing an occultation survey! We therefore assume that, for a given asteroid, the three most precise occultation measurements will be used to constrain the initial state, and information on perturbations will begin to accrue only from the fourth most accurate occultation observation. Thus, a target which does not produce ≥ 4 measurable occultations does not contribute at all to the σ_{tot} figure of merit. When forecasting the actual number of occultations, we assume each target draws from a Poisson distribution with

mean given by Equation (3), rather than assigning every asteroid the mean number of observed occultations.

Equation (3) shows that $\langle N_e \rangle \propto a_T^{-2}$, thus the Trojans and particularly the TNOs will be starved for occultations, with $\langle N_e \rangle < 4$ unless we increase the array length well beyond 200 km. It is also clear that the number of observed events N_e scales with the number of telescopes N at a fixed choice of spacing, so our goal quantity will scale as $\sigma_{\text{tot}} \propto N^{-1/2}$, and in fact will scale more strongly than this because of the threshold $N_e \geq 4$ required to make a test particle useful.

3.4. Results: MBAs

Figure 3 shows, for one scenario of an MBA survey, the results of the above assumptions on a grid of MBA diameter versus Gaia source magnitude. The left panel shows the expected σ_x per occultation for the (d_T, m_G) combination. The right panel weights the σ_x^{-2} information per occultation in each region in $d_T - m_G$ space by the number of asteroids and the number of usable occultations per asteroid, to give the total information density in $d_T - m_G$ space, following Equation (2). An occultation is assumed to be useful only if ≥ 3 observations with lower σ_x have been accrued for that object.

The projected value of σ_{tot} in this scenario is just 4.4 cm, similar to the accuracy available from Mars ranging, with a nominal \$15,000,000 capital cost. This is quite encouraging, since the tidal field of a $5M_\oplus$ Planet X at 800 au will cause a

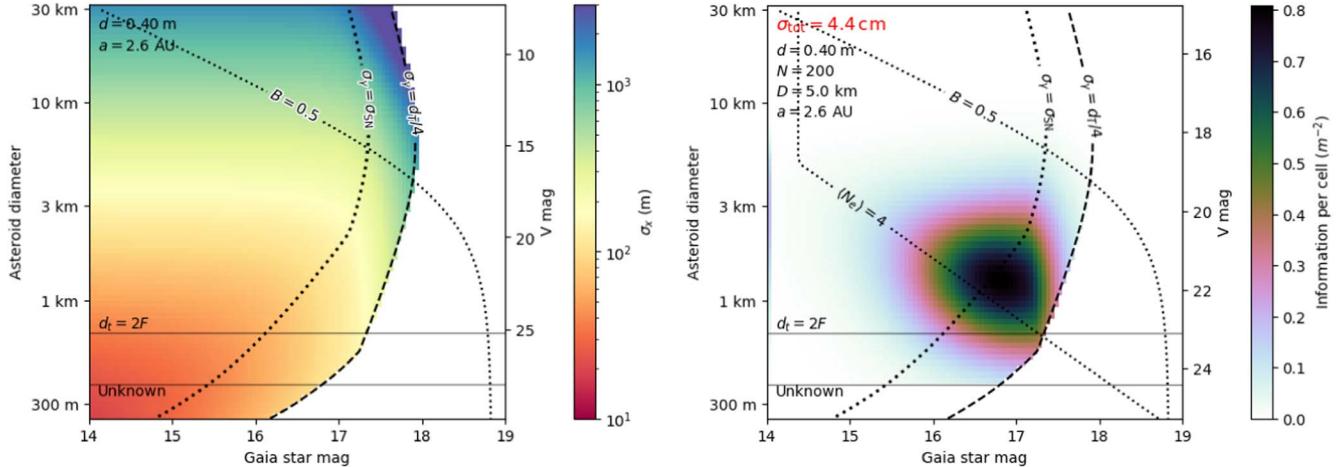


Figure 3. These figures illustrate the performance predicted by the parametric model for an occultation array targeting MBAs using a configuration of $N = 200$ telescopes of diameter $d = 0.4$ m at $D = 2$ km spacing, near the optimum for a \$15,000,000 nominal capital cost. The uncertainty on the rms size of an anomalous gravitational displacement after a 10 yr survey is just $\sigma_{tot} = 4.4$ cm. At left is the total measurement uncertainty in the occulter position for single-chord observations vs. the occulter diameter and magnitude of the occulted star. Overlaid are several important dividing lines; asteroids below the $d_f = 2F$ line have radius below the Fresnel scale, only partially blocking stellar light. Those below the “Unknown” line are too faint to be found in the nominal LSST single-epoch data. Those to the right of the $B = 0.5$ contour are substantially degraded by “background” light from the night sky and the asteroid’s reflected light. Those to the left of the $\sigma_\gamma = \sigma_{SN}$ contour are limited by shape noise in determining the CM of the asteroid. Those to the right of the $\sigma_\gamma = d_f/4$ curve have too much photon noise for the occultation to be securely detected. The right-hand plot shows the density of information on gravitational perturbations in the d_f-m_G plane, which is heavily concentrated in the region with the most numerous occultations—the asteroids near the Fresnel scale, with the faintest stars that produce detectable occultations. The $\langle N_e \rangle = 4$ contour marks the magnitude of the star that will on average be the fourth-brightest occulted by an asteroid of a given size. Occulters with $N_e < 4$ produce negligible information on gravitational perturbations.

shift of $\approx 5e_T$ meters from apsidal precession over 10 yr (where e_T is the occulter’s eccentricity) and a quadrupolar position oscillation with amplitude $\approx \pm 10$ cm on each orbit.

These plots show how the region of useful information is bounded at the faint- m_G end by having sufficient statistics to detect the occultation, ($\sigma_\gamma < d_f/4$). This bound moves to fainter m_G with larger telescope diameters d , and is independent of the number of stations N . The high-information region is bounded at the bright- m_G end by the need to accrue enough occultations over 10 yr to determine the state vector and begin to extract information on perturbations ($N_e > 4$). This bound improves linearly with N and with the survey duration T , and is largely independent of d . As the target diameter d_f decreases, the number of available occulters and available information increase rapidly, until these two constraints converge, which by coincidence is near the Fresnel scale in this scenario. The majority of information is arriving from asteroids near the $\langle N_e \rangle = 4$ line, with some targets below the line still providing their contribution if they end up with $N_e \geq 4$ occultations because of Poisson fluctuations in the number of occultations over the survey. The bulk of the information on gravitational perturbations comes from MBAs slightly above the Fresnel scale, since the number of such bodies increases rapidly with smaller diameter, and likewise is concentrated in the faintest source-star magnitudes whose occultations are detectable.

The parametric treatment of this section allows rapid exploration of the space of survey characteristics. We observe that configurations optimized under a cost constraint have the following characteristics:

1. The trade-off between telescope diameter d_f and number of telescopes N at fixed cost minimizes σ_{tot} when the cost of the telescope $C_t(d_f/0.5 \text{ m})^2$ is 10%–20% higher than the fixed station costs C_s . In other words, one should spread the funds over as many (smaller) telescopes as possible until the per-station costs are about to dominate.

When $C_s = \$40,000$, this leads to an optimal telescope size $d_f \approx 0.45$ m.

2. The constraining power of the array grows rapidly with the total cost cap, $\sigma_{tot} \propto \text{cost}^{-1.3}$. Higher capital funds allow larger N , which benefits the performance in two ways: first, the total number of observable occultations increases linearly with N ; second, the fourth-brightest occultation moves to brighter stars, which means that lower- σ_x measurements can now be applied to gravitational constraints.
3. Increasing the survey duration T increases the constraining power at an even higher exponent, since increasing T and increasing N reduce the σ_{tot} at the same rate—but increasing T gives gravitational anomalies more time to act, raising their resultant displacements.
4. Reducing the fixed costs per station is highly beneficial, roughly as $\sigma_{tot} \propto C_s^{-0.5}$, at fixed total costs since this moves the optimization toward more, smaller telescopes, improving $\langle N_e \rangle$.
5. The right panel of Figure 3 shows that LSST is capable of discovering nearly all of the MBAs that would make useful test particles.
6. The figure also shows most of the information is coming from occultations that are *not* strongly affected by background light. This is good news in that it makes the survey less sensitive to the degraded seeing and night-sky brightness that a geographically dispersed array might encounter.
7. It is also true that much of the MBA information is from objects close to the Fresnel size and “lucky” to obtain four measured transits. Both of these place stress on our simplified assumptions about the amount of information extractable from each occulter, and the latter might reduce our robustness against other nuisance effects such as nongravitational forces. The simulations in Section 4

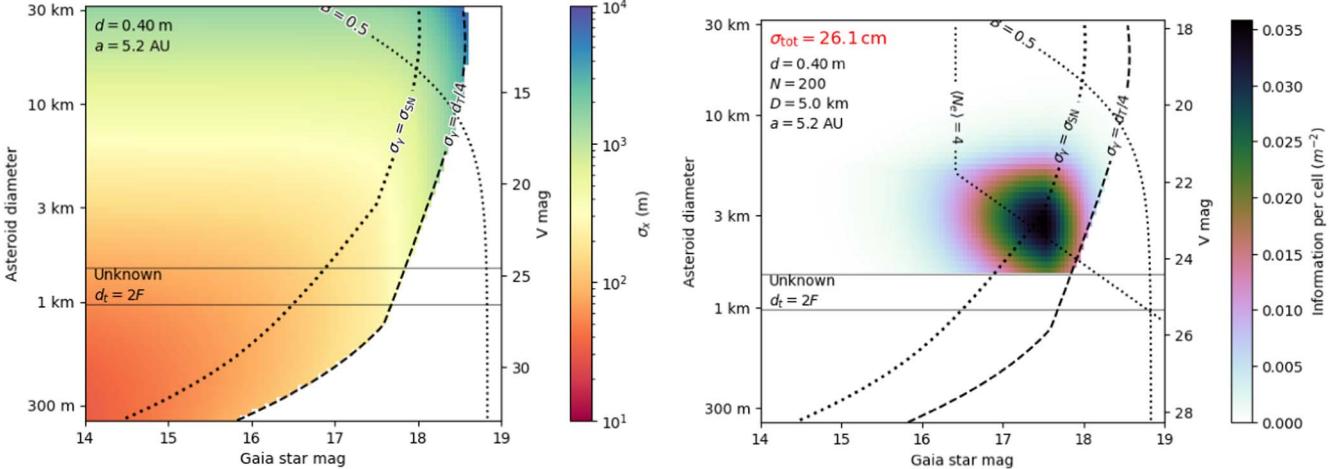


Figure 4. Summaries of the per-occultation measurement errors (left) and total information (right) for an occultation survey targeting Jovian Trojans. All curves have the same meaning as for the MBA case in Figure 3, except that the color scales are changed. The main configuration change is to move the telescopes further apart from each other to increase the occultation rates for the larger-diameter, sparser Trojan population. Analysis is in Section 3.5.

will test the robustness of the extracting of gravitational inferences from the positional data.

3.5. Results: Trojans

Using the Jovian Trojans as test bodies instead of MBAs changes the optimization of an occultation array in important ways, even though the $2\times$ increase in Earth–occultor distance changes the Fresnel scale by only $\sqrt{2}$. One issue is that, per Figure 1, there are $\approx 5\times$ fewer Trojans of a given size than MBAs. Furthermore, a Trojan of a given size is 3 mag fainter than its MBA counterpart, which means that LSST’s nominal survey is *not* capable of finding all Trojans above the Fresnel scale, as can be seen in Figures 4. Most importantly, the sky area subtended by the annual track of a Trojan is $4\times$ lower than for a similar-sized MBA, which means we must reoptimize the array to obtain $\langle N_e \rangle \geq 4$ for a substantial fraction of the observable Trojans.

We therefore investigate changing the telescope spacing from $D = 2$ km for the MBAs to 5 km for the Trojans, which will increase the number of observable occultations for Trojans with $d_T > 2$ km—which is, in fact, all of the LSST-discoverable sources. Figure 4 illustrates the result of a configuration near the new optimum for a \$15,000,000 nominal cost. The number and size of telescopes resemble the MBA optimum, but spread out over 1000 km. This wider-spaced array might be harder to maintain; it would, however, also be able to conduct the MBA survey, with the caveat that many forecasted MBA occultations would be “duds” that pass between stations, leading to a higher workload for the array.

The scaling relations of the optimal d_T and of σ_{tot} versus the cost, duration, and per-station costs remain essentially the same as the MBA case. Most of the occultations still occur in low-background situations. There are some practical differences, however:

1. Most of the information is now arising from occultations that are photon-noise limited rather than shape-noise limited.
2. The usable occulters are above the Fresnel size, so occultations will be deeper and will last longer.

3. The total number of Trojan occultations would be many times lower than the number of MBA occultations.

The reduced number of Trojan targets and the need to use larger occulters both act to reduce the statistical power of the Trojan survey relative to the MBA. Our \$15,000,000 scenario forecasts $\sigma_{\text{tot}} = 23$ cm, which is $\approx 5\times$ less sensitive to displacements than the MBA survey. It is also true, however, that the apsidal precession from a Planet X tidal force over a fixed time T scales as $a_T^{5/2}$ and the quadrupolar oscillation amplitude scales as a_T^4 , factors of 5.6 and 16 larger for Trojans than MBAs, respectively. Hence, our parametric investigations suggest similar constraining powers for studies of the two populations—and considerable gain if we can do most of both surveys with one array.

3.6. Results: TNOs

An occultation survey of TNOs for the purposes of constraining tidal forces would face more extreme versions of the challenges of a Trojan survey. First, the number of targets discovered and with good orbit determinations from LSST is very low, with a recent estimate of 30,000 (J. Kurlander 2024, private communication), with the smallest of them near $d_T = 100$ km. Thus, the shape-noise limit alone will degrade per-event accuracy to well above the Fresnel scale. The good news is that the longer, deeper occultations will allow detection of occultations of fainter Gaia stars. This increase in event rate will, however, be overwhelmed by the huge decrease in the area subtended by the apparent motion of the occulters. Anticipating that we will need maximal array cross section to obtain $N_e \geq 4$ occultations on the scarce targets, we consider a global-scale array with 100 telescopes at $D = 100$ km.

Another important distinction between the TNOs and the Trojans/MBAs is that the survey duration T is a small fraction of the orbital period. The TNOs are essentially in inertial motion and the deflection due to any tidal force will grow quadratically in time, whereas in the rest frames of the MBAs and Trojans, the applied tidal force has been rotating. A Planet X at 800 au will cause a deflection of ≈ 180 km over 10 yr, far larger than the signals applied to the MBAs.

Figures 5 diagnoses a scenario where $d_T = 0.4$ m telescopes (again, nearly optimal) are arrayed at 100 km intervals,

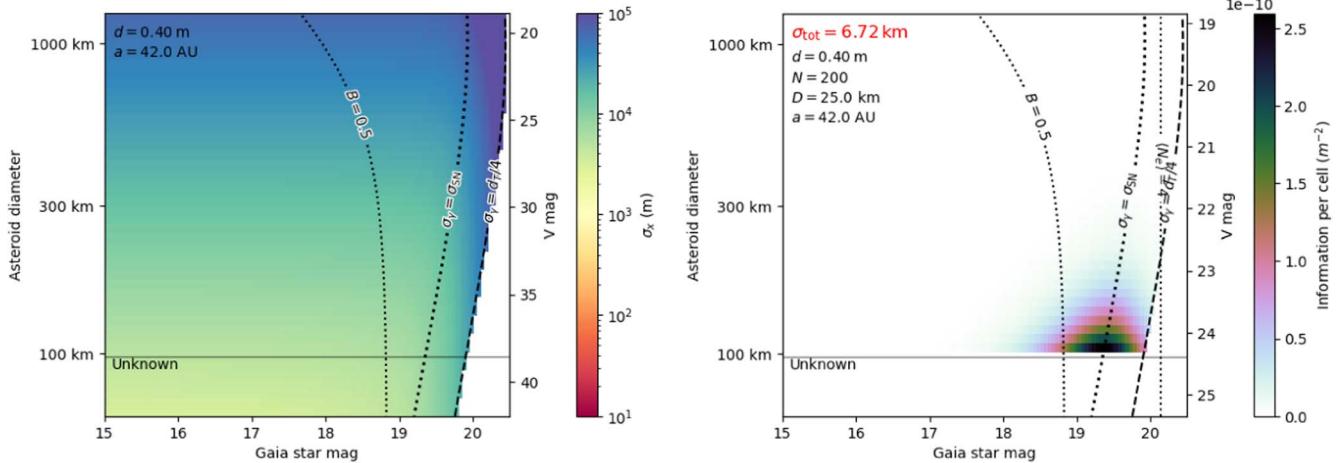


Figure 5. Summaries of the per-occultation measurement errors (left) and total information (right) for an occultation survey targeting TNOs. All curves have the same meaning as for the MBA case in Figure 3, except that the color scales and axis limits are changed. These plots assume $N = 100$ telescopes spaced $D = 100$ km apart, a much larger array than considered for MBAs and TNOs in order to obtain sufficient occultations per TNO. Further commentary is in Section 3.6.

spanning 10,000 km—a global-scale effort. At this point we can see that the TNOs attain $\langle N_e \rangle = 4$ events for $m_G \approx 19$. Most of the information is coming from fainter Gaia stars, and background counts produce most of the noise. The array does, however, attain $\sigma_{\text{tot}} = 1.3$ km, which is physically much larger than the MBA and Trojan precision levels. The much larger tidal effects on the TNOs, however, potentially allow for much stronger constraints on unknown masses.

In some sense, a TNO array would be technically easier, in that the events are much slower and rarer. Our cost formula would describe this as an \$8,000,000 array, but certainly operational costs would become much larger for a globe-spanning array than for an array that can be reached by truck on day trips from one or two maintenance facilities. The RECON network (M. W. Buie & J. M. Keller 2016) has deployed 64 $d_T = 28$ cm telescopes spanning a ≈ 2000 km north-south path from Kelowna, British Columbia in Canada to Yuma, Arizona in the USA. The RECON stations are operated by local volunteers on an occasional basis, with the primary goal of measuring shapes and sizes of TNOs and Trojans. While these are not full-time, autonomous stations such as we propose, RECON is a useful demonstration that useful sites at 50–100 km separation can be secured for telescopes.

4. Simulations

Given that the previous section's rough calculations of the sizes of gravitational perturbations measurable from a feasible occultation array suggest that secure detections of the tidal forces of an $M_X = 5M_{\oplus}$, 800 au planet should be possible, we conduct more detailed simulations of such an experiment. Rather than use σ_{tot} as a measure of effectiveness, the simulations use a Fisher-matrix method to estimate the uncertainties in the tidal field itself, after marginalization over all other unknown parameters of solar system dynamics, including the initial state vectors of all test bodies and gravitating bodies. Using full distributions of MBA and Trojan orbits and sizes, and the real distribution of Gaia stars on the sky, the simulation realizes a full population of occultation events and calculates the positional uncertainties—both along-track and cross-track—expected to be obtained for each event given an observing scenario.

This methodology enables more accurate treatment than the parametric models for many effects, starting with fully realistic distributions of test-body orbits and sizes, the true dynamical perturbations a tidal force would create, and degeneracies of tidal forces with other parameters of the solar system model and nongravitational forces. Other effects treated more accurately here include degradation of measurement errors with air mass and moonlight; observability of events given twilight times, weather fractions, and air mass limits; gains in accuracy from multichord occultations; and the uneven distribution of Gaia stars on the sky.

We do not include TNOs in this analysis due to the global scale required for a TNO survey to be used for measuring tidal forces. We do, however, predict how many TNO occultation events would be measurable from a MBA/Trojan-focused array. These observations, though insufficient to provide constraints on Planet X, would increase the number of detailed TNO shapes, sizes, and albedos known to us by $>100\times$.

4.1. Event Simulation

Our survey simulation generates a random sample of occultation events based on the stellar density throughout the path of the target asteroids. We compute a pixelized sky map of the local stellar density, in bins of stellar color and magnitude, from the Gaia Data Release 3 catalog. This allows us to quantify the expected events and their uncertainties without the precise knowledge of minor-body orbits and telescope locations that would be needed to predict each and every event.

A random subset of $\mathcal{O}(10^3)$ orbits is selected from those available in the Minor Planet Center database for each population—MBAs and Jovian Trojans. The full population of occulters is generated by drawing an absolute magnitude from the expected magnitude distribution of the population within the LSST discoverability limit, following estimates by LSST Science Collaboration et al. (2009). A diameter is then computed from the absolute magnitude and albedo through the known expression (A. W. Harris & A. W. Harris 1997)

$$\log d_T = 3.1236 - 0.5 \log q_T - 0.2H. \quad (3)$$

To capture the known bimodality in the albedo distribution of MBAs, we randomly select albedo values from the distribution

presented by Z. Murray (2023). For the Trojans, we assign $q_T = 0.12$ to the entire population, as this is the average for small Jovian Trojans reported by Y. R. Fernández et al. (2009).

Each simulated target asteroid is then assigned one of the random MPC orbits. For each 24 hr period of the simulated survey, the asteroid's position in the pixelized Gaia star map is found, and an occultation rate n_{occ} is assigned according to the length of the asteroid arc $\Delta\alpha$ and local stellar density n_G as

$$n_{\text{occ}} = n_G \frac{ND}{D_T} \Delta\alpha, \quad (4)$$

where D_T is the Earth–asteroid distance, and angles are expressed in radians. Occultations are drawn from Poisson distributions around the computed rate n_{occ} .

Generated occultations are then assigned a random time stamp within the given 24 hr period, and unobservable events are filtered out. For this purpose, the latitude is assumed to be 35°N, a rough estimate if the array were to be built in the southwest United States. Events are limited to night time by selection of those which occur when the Sun is lower than -18° altitude. Furthermore, events that happen at air mass $>X_{\text{max}} = 2.3$ are discarded, as well as those closer than 30° from a more-than-half illuminated Moon. Finally, we discard a random $\approx 30\%$ of the nights as having high cloud coverage.

4.2. Sources of Error

The expected astrometric error for a single occultation event is a combination of several uncertainties that require careful consideration. Already mentioned in Section 3.2 are the shape noise (σ_{SN}), photon noise (σ_γ), and Gaia astrometric uncertainties σ_G . While the latter is fixed for any given star, the two former ones require further modeling when performing the simulation. Additional sources of error include telescope location and timing of events, as well as angular sizes of stars, since they are assumed as point sources when computing photon-noise estimates.

4.2.1. Shape Noise

In order to estimate the magnitude of shape-noise errors, we simulate occultation data for a population of randomly oriented objects created from shapes found in the 3D Asteroid Catalog.⁵ We then attempt to predict the CM of each object, in the direction perpendicular to the array, as the average of occultation chord centers weighted by the length of the chords. For single-chord events, the rms error is $\sigma_{\text{SN}} = 0.1(d_T/2)$; for two or more chords, it approaches $\sigma_{\text{SN}} = 0.04(d_T/2)$.

In the direction of the array, assumed to be north–south, the CM location is estimated from the fixed locations of the telescopes involved in the event, being limited by the distance between telescopes or object diameter (whichever is smaller). The rms is therefore expected to be $\sigma \approx \min(d_T, D)/\sqrt{12}$.

Additional noise will be generated if a binary system is mistakenly assumed to be a single object, because only one member of the system is detected in occultation. In this case, our best CM prediction will be offset by a factor dependent on the binary mass ratio and separation. Near-Earth orbit observations indicate a binary fraction of 15% (P. Pravec et al. 2006). Assuming a similar figure for the MBAs, we look at the distribution of mass ratios and separations of known

kilometer-sized MBA binaries, and we find that most have CM offsets smaller or comparable to our nominal shape-noise error. We therefore do not model this additional error, and introduce the following caveat to our results: ideally, the survey should extend until enough detections of each target are made that additional scatter caused by a binary companion can be statistically inferred.

4.2.2. Flux and Photon Noise

Photon-noise error depends on the star's apparent magnitude, as well as background sky brightness. For each occultation event, star magnitude and color are randomly drawn from the color and magnitude distributions previously assigned to the event's sky location. Sky brightness as a function of Moon phase is estimated as specified in Appendix B. When the Moon is below the horizon, the dark-sky value is assumed. We further include a linear scaling of background flux with air mass X , as well as seeing proportional to $X^{3/5}$. Both relations are in agreement with seeing and sky brightness measurements by the Dark Energy Survey (E. J. Neilsen et al. 2019).

Equipped with stellar magnitudes and sky background values, the fluxes before and during the occultation events are computed, and thus the photon-noise error for each chord is estimated. The number of chords is drawn based on its probability for each given asteroid size and assigned distance between telescopes. With this draw, occultations that slip between telescopes are also discarded (when zero chords are assigned to them). Having multiple chords reduces the expected photon noise, and the shape noise as described above, but uncertainty propagated from the Gaia star's positional error is independent of the number of chords.

4.2.3. Telescope Location and Timing

Regarding the required accuracy in the timing of events, a clock error of 1 ms induces a ≈ 30 m error in Earth's position, and hence that of the asteroid, which is below the estimated shape noise for our target populations. Similarly, it is very easy to conduct geodesy to map each telescope's location on Earth to a level far below our estimated single-event uncertainty σ_x . The more difficult challenge, however, may be to limit systematic errors (i.e., correlated among all asteroids) in timing and geodesy so that they are below the combined uncertainty for all our survey data of $\sigma_{\text{tot}} \sim 4$ cm found in Section 3.4. This requires that systematic timing errors be kept below 1 μ s, and that systematic positioning errors on the stations' mean position be at most at the centimeter level. Both should be achievable with GPS devices.

4.2.4. Stellar Angular Diameter

The finite stellar angular diameter is not taken into account in our modeling of diffraction. If the stellar diameter approaches or exceeds the Fresnel angular scale of the occulter, it will broaden the ingress/egress light curves, and can also decrease the shadow depth, so we must consider it as an additional source of measurement error. We estimate the magnitude of this effect by generating temperature estimates for occulted stars, as follows. For every generated event, a star magnitude and color is assigned from Gaia catalog distributions at the given sky location. A temperature value is assigned for each color bin, corresponding to the median temperature of Gaia stars within that bin which have a temperature estimate in the

⁵ <https://3d-asteroids.space/asteroids/>

catalog. We then use the magnitude and luminosity relations,

$$M = m - 5 \log(d/10 \text{ pc}) - A, \quad (5)$$

$$L = L_{\odot} 10^{-0.4(M+B_{\text{cor}}(T)-M_{\odot})}, \quad (6)$$

where A is the extinction coefficient and $B_{\text{cor}}(T)$ is the bolometric correction, and the Stefan–Boltzmann law,

$$L = \sigma T^4 \pi R^2, \quad (7)$$

to get the angular diameter:

$$\alpha = 2R/d. \quad (8)$$

Notice that the result does not depend on the distance (the dependence on d from Equation (5) is canceled in Equation (8)). Extinction values are estimated based on an average for stars on the given pixelized cell. We use the bolometric correction model for the Gaia G filter from R. Andrae et al. (2018) for $T < 8000$ K, and the LTE+NLTE (non-local thermodynamic equilibrium) statistical model from M. G. Pedersen et al. (2020) for $T > 10,000$ K. The $8000 \text{ K} < T < 10,000 \text{ K}$ gap was filled with a linear model that would keep $B(T)$ continuous at $T = 8000$ K and $T = 10,000$ K. For most occulted stars, the resultant angular diameter is smaller than astrometric errors from other sources, so stellar angular diameter is not believed to have a significant effect in our results. To be conservative, however, we add the stellar angular diameter in quadrature to our other measurement uncertainties when the star is large enough for this to be relevant.

4.3. Dynamical Model

We now shift our attention to the propagation of astrometric measurement errors into uncertainties on external tidal forces by creating the following dynamical model of the solar system. As gravitationally active bodies, we include the Sun, the eight known major planets, plus Pluto and the 343 largest MBAs considered as active bodies by W. M. Folkner et al. (2014). The major planets have state vectors and masses as free parameters. The Sun has a free parameter for its mass and its oblateness J_2 , but a fixed state vector, in order to remove the coordinate reference frame degeneracy from the system. Pluto and the active asteroids also have fixed state vectors and free masses, since their low masses make it unimportant to know exact state vectors. We also model the gravitational potential of the remaining MBAs as arising from two circular rings of mass at $r = 2.06$ au and $r = 3.27$ au, respectively, and the azimuthally symmetric portion of the Kuiper Belt with two more rings at $r = 39.5$ au and $r = 43$ au, respectively.

There are two relevant nongravitational forces that act significantly on test particles (the observed MBAs and Trojans): solar radiation pressure (SRP) and the Yarkovsky effect. The accelerations attributable to these are parallel and perpendicular, respectively, to the Sun–asteroid vector, lying in the orbital plane, and each is inversely proportional to the body diameter, such that⁶

$$a_{Y,\text{SRP}} = k_{Y,\text{SRP}} \frac{a_{\odot}}{d_T}, \quad (9)$$

where a_{\odot} is the solar gravitational acceleration on the body.

⁶ There is a less important component perpendicular to the orbital plane.

In our previous paper (D. C. H. Gomes et al. 2023), we modeled each of the k parameters as global scaling factors. We improve this model by allowing each asteroid to have a free Yarkovsky tangential acceleration parameter. This change is required because our predicted per-object CM uncertainties are smaller than the expected displacements from the Yarkovsky force, the latter being of the order of a few kilometers for most MBA targets during a 10 yr arc. Therefore, the variance in Yarkovsky displacements due to different asteroid shapes and albedos will exceed the other measurement errors. Adding individual Yarkovsky accelerations for each asteroid as constrainable parameters solves this issue, with the only downside that each orbit solution needs an extra data point before it starts being useful for tidal-field information: Four initial occultations are needed instead of three. On the positive side, our occultation survey would successfully detect this Yarkovsky acceleration parameter on most of our targets, an achievement that otherwise would require many decades of observations (D. Hung et al. 2023). For SRP, which has a smaller contribution to orbital displacement, we keep a single global parameter, k_{SRP} .

The final element of the dynamical model is a set of five point masses at 400 au and different fixed positions. D. C. H. Gomes et al. (2023) show that a Planet X tidal field at any location in the sky can be expressed as a linear combination of the tidal fields from these five putative masses. Hence, forecasting the covariance matrix of these five mass parameters is sufficient to place bounds on the mass of a Planet X in any chosen sky direction. The quadrupole moment of the Kuiper Belt mass distribution generates a gravitational field on the Trojans and MBAs that is indistinguishable from a Planet X tidal field. We therefore add to the 5×5 covariance matrix an additional covariance arising from the positions and mass uncertainties of the largest individual TNOs, as well as quadrupole fields from the remaining TNO population, as modeled by the CFEPS project (J. J. Kavelaars et al. 2009; J.-M. Petit et al. 2011; B. Gladman et al. 2012). This calculation is described in more detail on our previous work (D. C. H. Gomes et al. 2023). The resultant “Kuiper Belt object (KBO) floor” is dominated by the mass uncertainties of the largest moonless TNOs.

We implement the dynamical model by running the REBOUND N -body integrator package (H. Rein & D. Tamayo 2016) to get variational derivatives of the planets’ and test bodies’ positions with respect to our model parameters. We start from a base simulation with the Sun and eight planets, to which the massive MBAs are also added. Test particles are included at the initial position of the surveyed MBAs and/or Trojans. We compute variational derivatives, with respect to the active-body parameters and the state vectors of the passive bodies, for the whole time range encompassed by existing historical data on the planets (1965 forward) up to the end of our proposed occultation survey. This is done performing two simulations that start at MJD 60000.0, which integrate forward and backward in time, respectively. These simulations are integrated with the `ias15` adaptive integrator, parallelized for 10 subsets of passive bodies and 30 subsets of MBAs. The small inconsistency between orbits predicted by each simulation, due to the presence of a different subset of massive MBAs, does not significantly affect our forecasts, since they depend exclusively on the variational derivatives, not on the state vectors.

The remaining model elements—ring potentials, solar J_2 , the mass of Pluto, and tangential acceleration due to the Yarkovsky effect—are added one by one to the simulations with two slightly different parameter values, so that their derivatives can be obtained through finite differences. To implement the ring potentials and nongravitational forces, we make use of the REBOUNDx package (D. Tamayo et al. 2019), as described in D. C. H. Gomes et al. (2023). Derivatives with respect to the SRP parameter are obtained by rescaling the derivatives with respect to the solar mass.

4.4. Fisher Uncertainties

Once equipped with our list of simulated observations and their uncertainties, as well as derivatives of model parameters with respect to the observables, we compute the Fisher matrix to estimate lower bounds on the uncertainties σ_M in the mass of a hypothetical Planet X at 4000 au. We follow the method delineated in D. C. H. Gomes et al. (2023). First, a Fisher matrix is generated for the set of existing observations of the major planets with respect to our global model parameters, namely the masses and state vectors of the planets, masses of ≈ 300 asteroids, collective mass rings for smaller MBAs and TNOs, and the tidal forces of interest. The observations we use for this step are the same ones enumerated in the aforementioned paper, and they consist of astrometry and radio ranging measurements. Next, for each small-body test particle, the Fisher matrix is augmented with rows and columns for its initial state vector parameters (and, in our current case, a Yarkovsky effect strength), and then the Fisher information from its simulated observational data is added in. We can then marginalize the Fisher matrix over this body's parameters. The Fisher matrix thus remains fixed in size even as we accumulate constraints from millions of test bodies. Finally, we invert the total Fisher matrix to form a covariance matrix, and extract the 5×5 submatrix corresponding to the five putative distant point masses. To the resultant covariance matrix we add the previously mentioned Kuiper Belt floor covariance. Then, for any given position in the sky, a transformation of the covariance matrix, as detailed in Appendix A of D. C. H. Gomes et al. (2023), yields us the lower-bound uncertainty σ_M on the mass of a Planet X in any chosen direction from the Sun.

We produce three distinct forecasts: first using only MBAs as test particles, then using only Trojans, and then for the scenario of a joint survey of both populations.

5. Simulated Constraints

We consider a grid of possible telescope arrays that vary between $N = 50$ and $N = 400$ telescopes, as well as telescope diameters between $d = 20$ cm and $d = 70$ cm. For each array configuration, we measure the constraining ability of the array by predicting the σ_M of Planet X on a full sky grid of possible point-mass positions, and selecting the 90th percentile value, i.e., the σ_M that can be achieved or surpassed over 90% of the sky in that particular survey configuration.

Three scenarios are presented in Figure 6. The first one (top left) assumes the array would be exclusively devoted to MBA occultations. In this case, the distance between telescopes is set at $D = 2$ km. There is no change in σ_M at larger D , since the majority of occulters are smaller than 2 km and the number of observed occultations becomes independent of D . The second

scenario (top right) assumes an exclusive Jovian Trojan survey, and sets the distance between telescopes to $D = 5$ km. Larger D is favored for the Trojan survey, since they are twice as far as MBAs, such that the smallest ones detectable by LSST are larger than for MBAs. The third scenario (bottom) assumes a joint MBA/Trojan survey, with $D = 5$ km again since the Trojans benefit and the MBA observations are not harmed by increased D .

The cost of each array configuration is computed according to Equation (1), so that the optimal configuration for any given cost can be easily determined. In all scenarios the optimal array diameter varies within the 30–40 cm range, and the optimal number of telescopes increases with the allowed budget. A 5σ detection of a $5M_\oplus$ body at 800 au, which is the smallest and most distant hypothesized Planet X invoked to explain alignment anomalies of distant TNOs, is achieved if $\sigma_M = 0.125M_\oplus$ (the tidal field of such a planet would be 8 times weaker than an equal mass at 400 au). Figure 6 shows us that a cost-optimized array capable of such detection would be slightly above \$15,000,000 for an MBA-exclusive survey. Trojans would not independently allow such a Planet X constraint, but they do slightly help the MBA result in a joint survey, making the same precision cross below \$15,000,000. Optimal configurations in both these scenarios are around $N \sim 200$ telescopes and $d \sim 40$ cm. Figure 7 shows the forecasted uncertainty σ_M of Planet X as a function of its location in the sky for a joint $N = 200$, $d = 40$ cm scenario.

To assess the feasibility of performing observations at or close to the predicted event frequency, we assume that a single measurement uses ~ 2 minutes of telescope time. The optimal configuration is therefore limited at ~ 6000 events \times telescopes per hour. Our simulations recorded hourly rates of successful observations (on nights with clear sky) varying in the 2000–3000 range. This means that, if exactly one telescope and time slot is required for each of 3000 successful events, the “subscription factor” would be $f \approx 0.5$ on the busiest nights, and if the events are randomly distributed over stations and time, a fraction $e^{-f} - (1 - f) \approx 0.1$ would overlap another event and be lost. If, however, two or more time slots are needed for each success, either because the orbit uncertainty spans multiple telescopes or the targets are small enough to slip between stations, then $f \approx 1$ or more, meaning $\approx 40\%$ or more of events would be lost to resource conflicts. This is most concerning at the initial stage of the survey, when target ephemerides are limited by LSST astrometry, and we expect a few failures before attaining the first successful observation of each target. Therefore, we would definitely benefit from array design options that push toward higher event frequency capabilities, and/or faster event cycles than 2 minutes. If the survey, before most objects get their first detection, is ultimately limited by the event load, we might need to extend survey duration to recover our predicted constraints.

In Figure 8, we see the decrease of σ_M for the joint survey as a function of the number of telescopes for both the $d = 30$ cm and $d = 40$ cm scenarios. Additional dotted lines show the constraints that would be attainable were there no degeneracy with Kuiper Belt tidal fields, i.e., computed out of the original covariance matrix, without a KBO floor addition. We find a strong gain in sensitivity for additional investment in N , with $\sigma_M \propto 1/N$ approximately—much better than the typical square-root relation. With the assumed Kuiper Belt floor, however, there is little gain in σ_M for $N \gtrsim 200$ and the associated cost

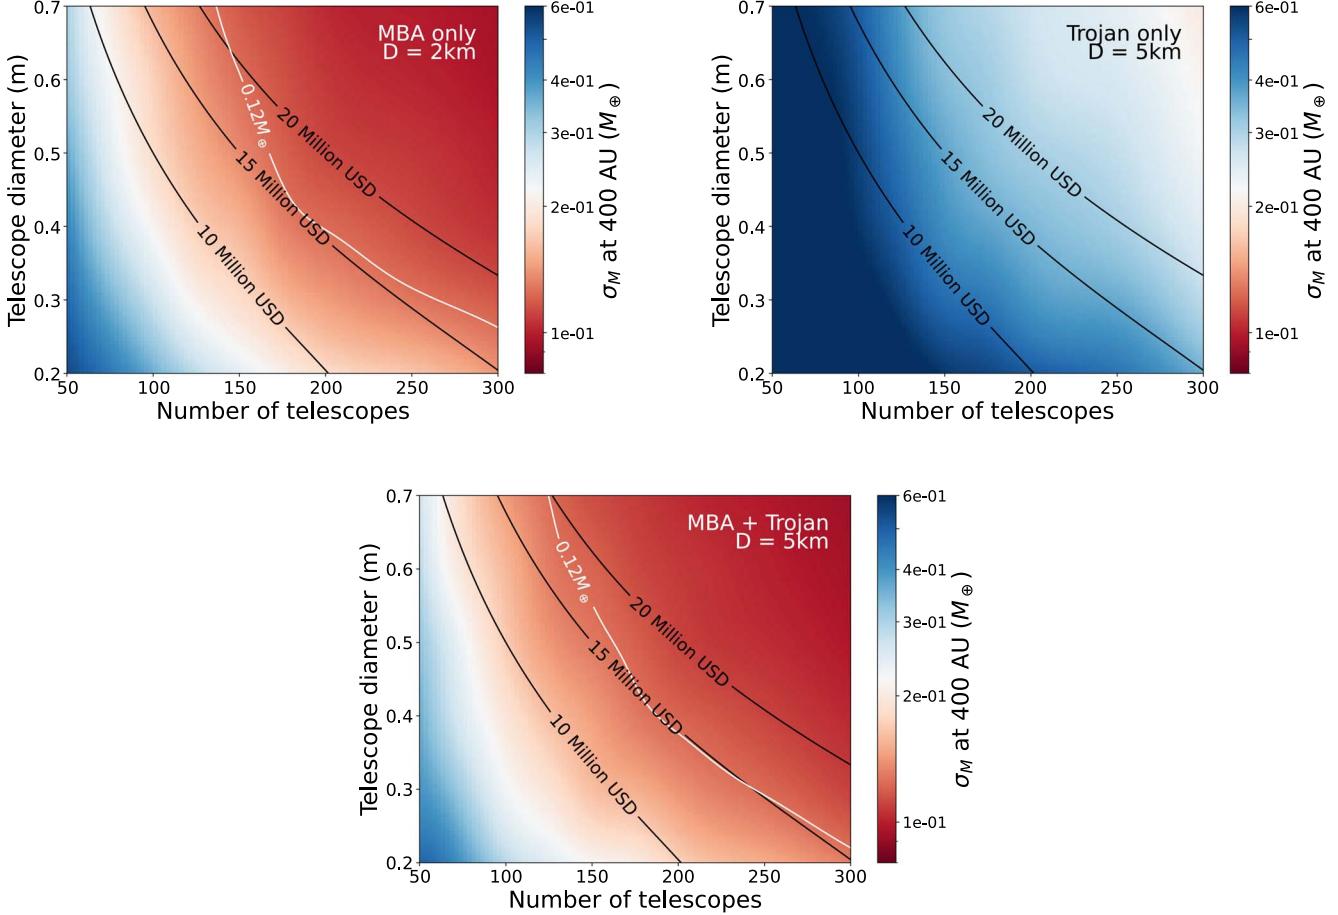


Figure 6. Value of σ_M under which a point source at 400 au would be constrained in 90% of the sky. Black contours are lines of equal cost. The white contour at $\sigma_M = 0.125 M_{\oplus}$ corresponds to a 5σ detection at 800 au. Top left: results for MBA survey with 2 km distance between telescopes. Top right: Jovian Trojan survey with 5 km distance between telescopes. Bottom: combined MBA and Trojan survey with 5 km distance between telescopes.

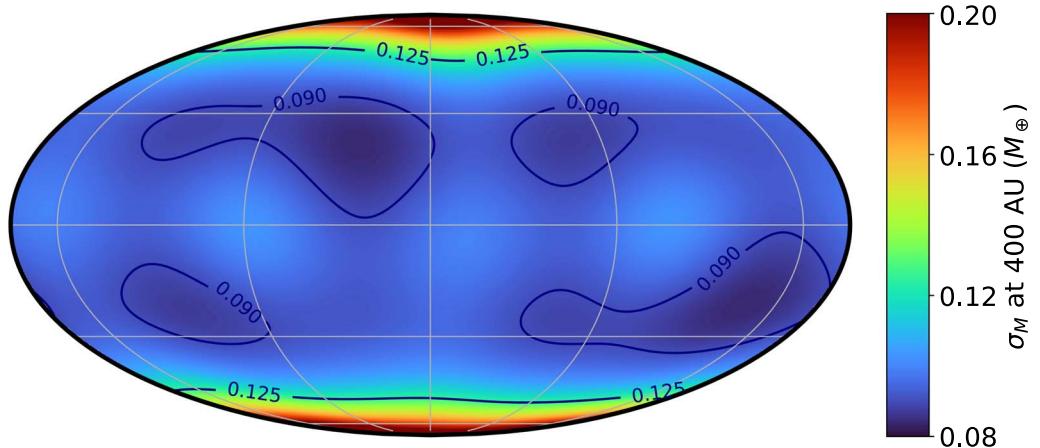


Figure 7. Lower bound on the uncertainty σ_M for a point source at 400 au as a function of its location in the sky. We use the Mollweide projection and ecliptic coordinates with the vernal equinox at center.

increase—although one would be measuring the KBO mass distribution, and other science benefits of the project beyond measuring σ_M would improve. Finally, Figure 9 shows the cumulative fraction of the sky where the uncertainty σ_M falls below a certain level, in two different scenarios: using exclusively historical data, and adding occultations assuming the optimal array to historical data.

We now consider the optimal $\sim \$15,000,000$ array ($N \sim 200$, $d \sim 40$ cm) and ask how effective such a configuration would be toward detecting TNO occultations. We use sample orbits from the CFEPS L7 population (J. J. Kavelaars et al. 2009; J.-M. Petit et al. 2011; B. Gladman et al. 2012) and find that, for a 10 yr survey, the array measures one occultation by 6% of the simulated objects (multiple

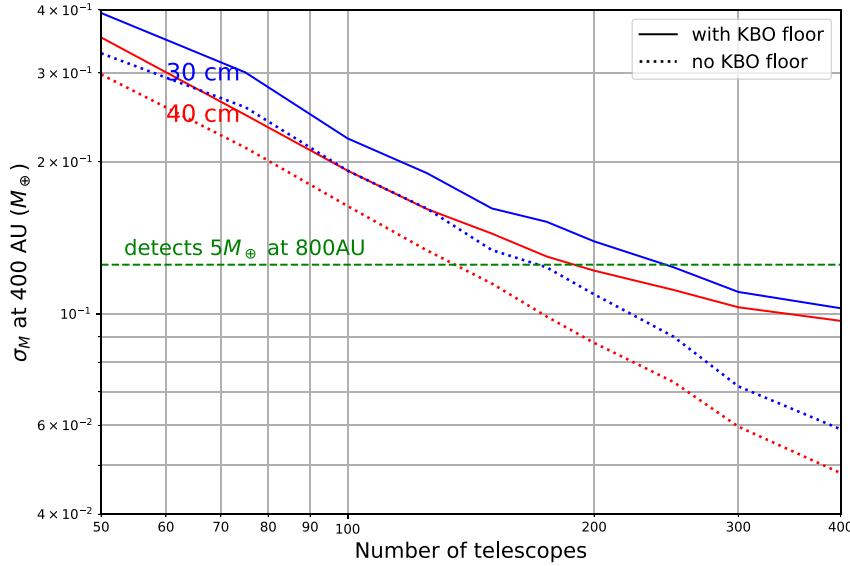


Figure 8. Value of σ_M for a point source at 400 au that would be achieved or surpassed for 90% of possible directions, for a joint MBA/Trojan survey. Blue line assumes 30 cm telescopes; red line assumes 40 cm. Solid lines include fundamental degeneracy of the tidal signal with the uncertainties of large TNO masses and asymmetries in the Kuiper Belt (KBO floor). Dotted lines show the predicted uncertainty disregarding this limitation.

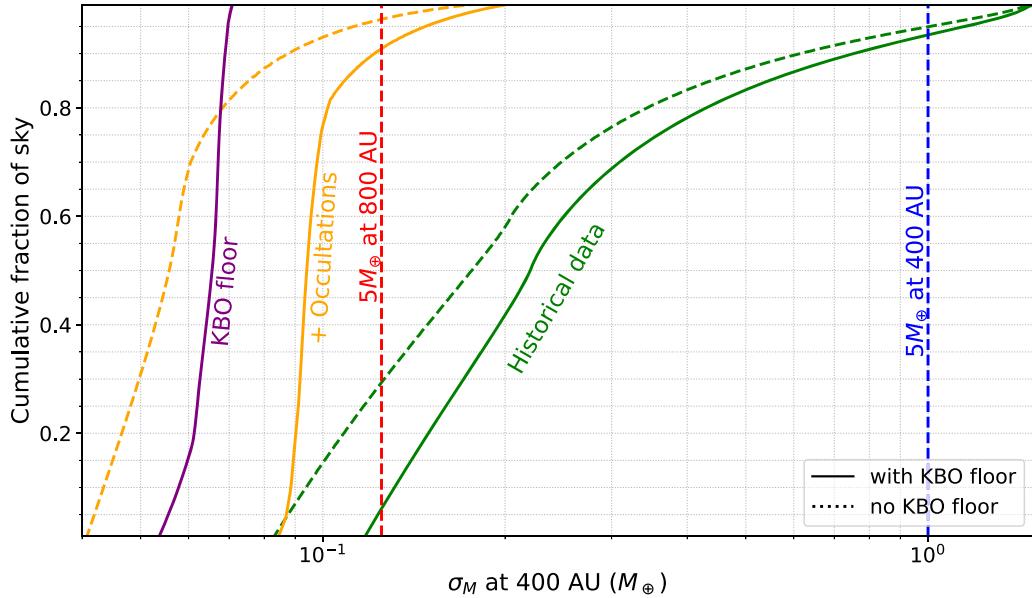


Figure 9. Fraction of the sky where the lower bound on σ_M lies below a given value. Dashed lines disregard the degeneracy of a Planet X tidal field with Kuiper Belt tidal-field uncertainties.

occultations for a TNO are rare). For the assumed number of 30,000 LSST-detectable TNOs, a 6% rate would mean 1800 distinct TNO shapes. Though such data will not provide constraints on orbital perturbations, a single occultation is still useful. A single occultation of these $d_T \gtrsim 100$ km objects with chords recorded at 5 km intervals will return a detailed projected size and shape measurement of the TNO, leading to a better understanding of the size, albedo, and contact-binary rate for a substantial number of TNOs in all dynamical classes. It will also yield more accurate predictions of subsequent occultation events, which could potentially be observed by other occultation campaigns.

6. Conclusions

We present the expected outcome of a proposed dedicated occultation array with $\mathcal{O}(100)$ telescopes, cost-optimized to constrain minor planets' orbits for detection of yet undiscovered tidal fields. An initial parametric analysis shows comparable results for arrays targeting MBAs and Jovian Trojans, with the $5\times$ weaker spatial precision for Trojans compensated by the presence of larger tidal perturbations on their orbits. Combination of data from all targets would provide centimeter-level sensitivity to tidal perturbations. A TNO-focused array would require different logistics, due to the slower apparent path of these objects on the sky and smaller known population.

To achieve the three required observations that initially constrain the orbit, such an array would need to be spread out across a $\sim 10,000$ km baseline between telescopes, turning the project into a global effort.

Survey simulations provide further insight into the MBA and Trojan array outcomes, by considering stochastic factors that affect the observability of each independent event. The simulations yield Fisher-matrix estimates of the σ_M of the mass of any undiscovered Planet X lurking in the solar system, including the effects of degeneracies with other sources of orbital perturbation. We find that MBA surveys achieve $2\text{--}3\times$ lower uncertainties than the Trojan ones, which is not as optimistic for the Trojans as our parametric analysis. We suspect that this results from the lower typical number of occultations observed for an individual Trojan than for an MBA, which makes it more difficult to disambiguate a tidal force from other dynamical perturbations. We consider which scenarios could yield a 5σ detection of a $5M_{\oplus}$ Planet X at 800 au on at least 90% of the sky. The lowest-cost option is a joint MBA/Trojan survey with $N \approx 200$ telescopes spaced at $D = 5$ km (i.e., a 1000 km array), each with diameter $d \sim 40$ cm, and a very crude construction cost estimate of $C \approx \$15,000,000$.

When compared with the previous suggestion of constraining a Planet X signal with LSST astrometry of Jovian Trojans (D. C. H. Gomes et al. 2023), an occultation survey is much more promising, mainly because it gathers information from the faintest end of the asteroid size distribution. LSST astrometry is assumed, however, as a prerequisite for accurate occultation predictions.

The constraint on missing mass in the solar system would not be the only science benefit of the small-telescope linear array. Our survey would allow the study of other perturbations to asteroid orbits, constraining Yarkovsky acceleration parameters and detecting mutual deflections by close asteroid encounters. These data, therefore, would contribute to a better understanding of asteroid densities, albedos, and thermal properties. Such an array would also provide detailed shapes for ~ 1800 TNOs, characterizing the distribution of sizes, shapes, albedos, and contact-binary rates of TNOs in many different dynamical states. No other envisioned observation could yield such information for a large sample of pristine products of planetesimal formation.

Characterizing orbital perturbations with the proposed level of accuracy can lead to further discoveries in physics. The

hypothesis that primordial black holes might be the constituents of dark matter predicts rates of close encounters between such objects and solar system bodies. T. X. Tran et al. (2024) point out that resultant perturbations would carry the particular signals of high perturber velocity ($v \approx 200 \text{ km s}^{-1}$) and nonecliptic trajectory, and would be distinguishable from other interstellar perturbers in that baryonic objects of comparable mass would be easily identifiable. Their proposed method of detection requires high-precision astrometry of multiple solar system objects, so that perturber information can be inferred from correlation between perturbations on different orbits—this aligns exactly with the goals of our array.

Sensitivity to centimeter-scale perturbations out to 5 au in a joint MBA/Trojan survey would yield constraints on modifications to general relativity as well. Looking beyond occultation observations, an array of hundreds of agile robotic telescopes, with a total collecting area equivalent to a ≈ 6 m telescope, could conduct other investigations between scheduled occultation events, such as targeted or blank-sky monitoring for transients. This would include monitoring of sources that are too bright or fast for LSST or that require cadences that do not fit in the LSST survey plan.

As for the location of the array, we are limited by latitude, since a high-latitude survey would be blind to significant portions of the ecliptic plane below the horizon. We also require a north–south stretch of $\mathcal{O}(1000)$ km with good observing weather, low light pollution, accessible terrain, and, preferentially, some preexisting infrastructure, namely a road, allowing for easy deployment of telescope stations and subsequent access for repairs and maintenance, and—if no cellular network is in place—collection of physical data storage media. Five candidate locations, selected with these conditions in mind, are the southwest United States, northern Chile, north-central Australia, Namibia/South Africa, and Saudi Arabia.

Acknowledgments

This work was supported by National Science Foundation grant AST-2205808. We thank Marc Buie, Matt Lehner, and Zachary Murray for helpful discussions.

Appendix A Symbol Glossary

Table 1 displays the main symbols used throughout this paper, their meaning and default value, where applicable.

Table 1
Symbol Glossary

Symbol	Meaning	Default Value
Observation specifications		
d	Telescope diameter	...
D	Telescope array spacing	2 km
N	Number of telescopes in array	...
C_s	Fixed cost per observing station	\$40k
C_t	Cost per $d = 0.5$ m telescope, scaling with area	\$60k
T	Survey duration	10 yr
λ_0	Central wavelength of passband	600 nm
η	Quantum efficiency of atmosphere/telescope/detector	0.5
f_{duty}	Duty cycle (fraction of occultations observed)	0.21
X	Air mass of observation	...
X_{max}	Maximum observable air mass	2.3
β_{min}	Minimum observable solar elongation	50°
θ_{ap}	Angular radius of photometry aperture	1"
Occulter specifications		
a_T	Semimajor axis of occulter (target)	...
d_T	Diameter of occulter	...
q_T	Geometric albedo of occulter	0.10
D_T	Distance to occulter	a_T
v_{\perp}	Relative transverse velocity of observer and occulter	$v_{\oplus} = 3 \times 10^4 \text{ m s}^{-1}$
m_{lim}	Faint limit of known (LSST) occulters	$m_r = 24.4$
$N_T(>d_T)$, dN_T/dd_T	Cumulative, differential counts of occulters	...
Source-star properties		
$n_G(<m)$, dn_G/dm	Cumulative, differential sky densities of usable Gaia stars	...
Derived quantities		
F	Fresnel length $\sqrt{\lambda_0 D_T / 2}$...
t_F	Fresnel crossing time F/v_{\perp}	...
t_0	Time of chord center	...
τ	Characteristic timescale of occultation	t_F
$f(u)$	Occultation light curve as a function of scaled time $u = (t - t_0)/\tau$...
n_b	Rate of background photocarriers	...
n_{\star}	Unocculted rate of star photocarriers	...
B	Background ratio n_b/n_{\star}	...
m_G^F	G magnitude of star producing four counts per t_F	...
m_G^B	G magnitude of star for which $n_{\star} = n_b$...
ρ	Scaled occulter radius $d_T/2F$...
σ_t	Total uncertainty on t_0	...
σ_x	Total along-track uncertainty on chord center	$v_{\perp} \sigma_t$
σ_{γ}	Contribution to σ_x from photon noise	...
σ_{SN}	Contribution to σ_x from occulter shape	0.05 d_T
σ_G	Contribution to σ_x from Gaia position	DR5 spec
N_e	Number of observed occultations per asteroid	...

Appendix B

Measurement Errors on Chord Midpoints

A photon-counting aperture photometer will be drawing a Poisson sample of events from an occultation light curve with a time-dependent rate given by

$$n(t) = n_b + n_{\star} f\left(\frac{t - t_0}{\tau}\right), \quad (\text{B1})$$

where n_b and n_{\star} are the count rates generated by background sources and the unocculted star, respectively. In this context, “background” refers to any source other than the star being occulted, including diffuse zodiacal and atmospheric emission, light reflected from the occulting asteroid, or neighboring stars, all of which can be expected to remain constant over the fraction of a second when the telescope is in ingress or egress of the shadow.

The function f gives the fraction of stellar intensity making it past the occulter, i.e., the light curve relative to the time of passage t_0 of the geometric shadow edge over the telescope, with time specified in units of some timescale τ . What is σ_t , the expected rms uncertainty in t_0 given the data D , where D in this case is a list of arrival times of detected photons? There is not, to our knowledge, any general analytic form for this “change detection” problem. We construct an approximate formula for σ_t based on analytic results for some limiting cases, and results of maximum-likelihood estimators applied to simulated photon streams.

The Cramer–Rao theorem bounds σ_t using the Fisher information on t_0 , which in this case can be written as

$$\sigma_t^{-2} \leq \left\langle \left(\frac{\partial \log p(D|t_0)}{\partial t_0} \right)^2 \right\rangle$$

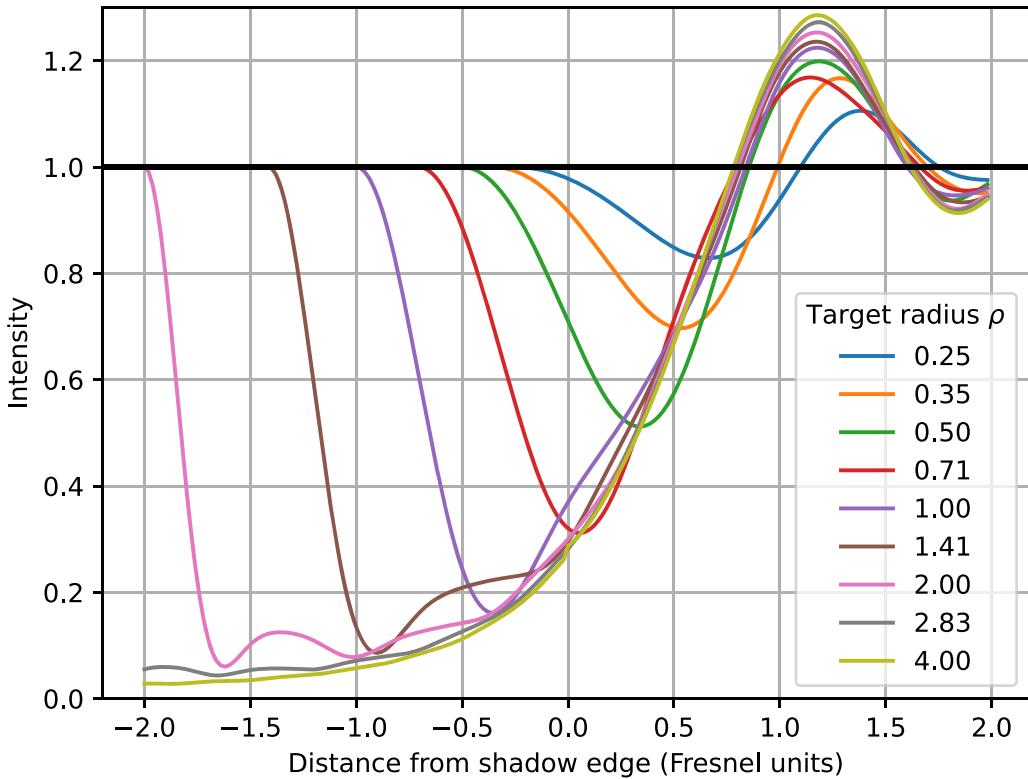


Figure 10. Light curves for the egress portion occultations of a point source by a spherical occulter of radius ρF , where $F = \sqrt{a\lambda/2}$ is the Fresnel scale. The horizontal axis can be considered either the distance from the geometric shadow edge (in units of F) or the time since the crossing of the geometric shadow edge (in units of $t_F = F/v_{\perp}$).

$$= \int dt \frac{[n_{\star}f'(t/\tau)/\tau]^2}{n_b + n_{\star}f(t/\tau)}$$

$$= \frac{n_{\star}\tau}{\tau^2} I(B),$$

$$I(B) \equiv \int du \frac{[f'(u)]^2}{B + f(u)}$$

$$\Rightarrow \sigma_t \geq \frac{\tau}{\sqrt{n_{\star}\tau}} I(B)^{-1/2},$$

where we have defined $B = n_b/n_{\star}$ as the background ratio. In the first equality, we have exploited the fact that the Fisher information in a Poisson distribution with a mean of $\lambda(q)$ for some parameter q is equal to $(d\lambda/dq)^2/\lambda$, and we can divide the light curve into statistically independent time segments of duration Δt with expected mean count rates of $\lambda = n(t)\Delta t$.

For the ingress or egress of occultation, the relevant timescale is $t_F = F/v_{\perp} \approx \sqrt{\frac{a\lambda}{2}}/v_{\oplus}$, the time for the Fresnel scale to cross a telescope moving at v_{\perp} across the line of sight to the target. Equation (B2) shows that the uncertainty on the start or end time of a chord is the Fresnel time divided by the square root of the number of source photons collected during the Fresnel time. The $I(B)$ factor encodes dependence on the size/shape of the target and the effect of background.

For a simple model in which $f(u)$ transitions linearly from 0 to 1 over a range $\Delta u = 1$, i.e., a transition time τ to complete shadow, the Cramer–Rao inequality yields

$$\sigma_t \geq \frac{\tau}{\sqrt{n_{\star}\tau}} [\log(1 + 1/B)]^{-1/2}. \quad (\text{B2})$$

The Cramer–Rao theorem does not prescribe an estimator for the time of ingress/egress, it just bounds any unbiased estimator's performance. We might expect the Fisher value to be attained when $\sigma_t \ll \tau$, when the differentials used in the Fisher matrix are applicable. We tested a straightforward estimator for the time of egress, namely the expectation value $\hat{t}_0 = \int dt_0 t_0 \mathcal{L}(D|t_0)/\int dt_0 \mathcal{L}(D|t_0)$. We find that equality in Equation (B2) is indeed attained, within $\approx 10\%$, by this estimator when $\sigma_t \lesssim 0.5$ and $B \gtrsim 0.02$. This gives us confidence to use the Fisher estimator for σ_t with the more complex $f(u)$ light curves obtained from diffraction, for cases when the uncertainty will be well below the Fresnel time.

To estimate typical $I(B)$ factors for diffractive occultations, we follow T. C. Nihei et al. (2007) by using the results of scalar diffraction theory for spherical bodies (G. E. Sommargren & H. J. Weaver 1990). Figure 10 plots the transmitted intensity function f for the egress portion of a central chord across a body of radius ρF . Each light curve is integrated over an octave of wavelength $2\lambda_0/3 < \lambda < 4\lambda_0/3$, and the horizontal axis is the variable $u = (t - t_0)/\tau$, with t_0 at the geometric shadow edge. Occulters of all size tend to follow a common curve in the neighborhood of the geometric shadow edge, with smaller bodies essentially having a maximum depth of shadow that truncates the curve. For each combination of ρ and background factor B we evaluate the integral in Equation (B2), using only

the time period between the minimum and maximum of the light curves in Figure 10. This period of time shows a largely monotonic rise, avoiding the central “Arago spot” and any oscillatory regions, which are likely to be strongly dependent on the detailed shape of the occulter.

The following functional form reproduces the resulting values of σ_t to within <5% over the relevant range of target radius ρ and all values of background ratio B :

$$\sigma_t^{\text{Fisher}} = 0.77 \frac{\tau}{\sqrt{n_* \tau}} \times [1 + (0.35/\rho)^{5/2}] \times \sqrt{1 + m_B(\rho)B},$$

$$m_B(\rho) = \min(1 + \rho^2, 1.8).$$

This result assumes central chords of the sphere. Numerical integration of the Fisher information for light curves for chords off-center by a fraction b of the occulter radius show weak degradation of σ_γ , by factors $\approx(1 - b^2)^{-1/8}$, to within 10% accuracy, which means that the mean information from random sampling of $b \in [-1, 1]$ is only $\approx 10\%$ lower than the calculation assuming central chords.

We need next to consider cases in which the Cramer–Rao lower bound is $\gtrsim \tau/2$. In such cases, occultations of duration $\lesssim 2\tau$ will be essentially undetectable, which means that this case is applicable only to sources with $\rho > 1$. For such objects, the shadow is dark, with $f(u)$ decreasing to $\lesssim 0.1$. Furthermore, if the transition from $f = 1$ to $f \approx 0$ is not resolved by the data, then the shape of this transition is unimportant, and we can assume an instantaneous drop. This leaves us in the simple geometric-optics limit, whereby the photon rate drops instantaneously from $n_* + n_b$ to n_b . Scaling arguments then tell us that we must have $\sigma_t = S(B)/n_*$ for some function S of the background-to-source flux ratio B . The case of $B = 0$ can be treated analytically: The estimator will clearly involve just the time of the first (last) arriving photon for egress (ingress), and one can show that $S(B = 0) = 1$. We can also expect $S(B) \propto B$ for $B \gg 1$, from the relation that the S/N of the dip in flux over a time period t is $\text{S/N} = n_* t / \sqrt{n_b t}$. We find an empirical approximation to $S(B)$ for geometric shadows that matches both expected limits, and is accurate to $\approx 10\%$ in between for simulations using our likelihood-based estimator for t_0 :

$$\sigma_t^{\text{geom}} = \frac{1}{n_*} \left[1 + \left(\frac{e}{\log(1 + 1/B)} \right)^2 \right]^{1/2}. \quad (\text{B3})$$

Our final estimator for the photon-noise uncertainty in the position of the center of the chord is to join σ_t^{Fisher} with σ_t^{geom} at the point where they cross, which is typically when $\sigma_t \approx \tau/2$. We increase the Cramer–Rao bound by a factor 1.1, and divide each result by $\sqrt{2}$ to account for the chord midpoint being the average of the ingress and egress times. Then, the time uncertainty is converted to a distance uncertainty with a factor of v_\perp :

$$\sigma_\gamma = \frac{v_\perp}{\sqrt{2}} \times \max$$

$$\times [1.1 \sigma_t^{\text{Fisher}}(\tau = F/v_\perp, \rho, n_*, B), \sigma_t^{\text{geom}}(n_*, B)]. \quad (\text{B4})$$

It is also the case that occultations for which $\sigma_\gamma \gtrsim d_T/2$ will not be reliably detected as decrements in stellar photon fluxes, and hence offer essentially zero information on the asteroid’s orbit.

Examining the above equations as a function of n_* , or equivalently the source star’s magnitude m_G , we find that there are two important inflections in the behavior of σ_γ at fixed values of n_b . One is the value n_*^F such that $n_*^F t_F = 4$, four photons being detected during the ingress/egress ramps. This is roughly the point at which the Fisher approximation begins to fail, and the Fresnel time is no longer resolved by the data. Assuming a one-octave filter centered at 600 nm, the magnitude in the Gaia G band at which this occurs is

$$m_G^F = 17.6 - 2.5 \log_{10} \left[\left(\frac{a}{2.6 \text{ au}} \right) \left(\frac{0.5 \text{ m}}{d} \right)^2 \left(\frac{0.5}{\eta} \right) \right], \quad (\text{B5})$$

with η being the QE of the atmosphere/telescope/detector combination. The other important inflection point is when the background rate equals the source rate, $B = 1$. This magnitude m_G^B is dependent only on the filter bandpass, sky characteristics at the site, and the radius θ_{ap} of the photometric aperture. We estimate this quantity using on-sky data for similar filters on the Dark Energy Camera on the Blanco telescope at Cerro Tololo⁷ and from the Megacam on the Canada–France–Hawaii Telescope.⁸ These estimates are in rough agreement with each other; we take the higher-background result as a conservative choice, keeping in mind that a real occultation array would need to be at lower altitude with higher background than these mountaintop observatories. For dark skies at zenith, a typical value for star–sky equality is

$$m_G^B = 19.6 - 5 \log_{10} \left(\frac{\theta_{\text{ap}}}{1''} \right). \quad (\text{B6})$$

This is faintward of m_G^F for MBA occulters and a 0.5 m telescope. Figure 2 plots the measurement error σ_γ of a chord midpoint position versus m_G for a nominal MBA with diameter well above $F \approx 300$ m, and a particular choice of telescope diameter d . For stars brighter than the inflection points, $\sigma_\gamma \propto f_*^{-1/2}$, but faintward of the inflections, stars lose power more rapidly, with $\sigma_\gamma \propto f_*^{-3/2}$. There is thus little to be gained by observing occultations of stars fainter than these inflection points.

We have not addressed the effect of scintillation on the chord detection. The timescale $t_F = 11 \text{ ms} \times (a/2.6 \text{ au})^{1/2}$ for the light curve is only slightly shorter than the scintillation timescale implied by a 10 m s^{-1} wind crossing a telescope of size $d = 0.4$ m, so there will be some scintillation effects on the relevant portion of the light curve. Figure 10, however, suggests that occultations for $\rho > 0.5$ have a depth greater than 50%, which is substantially larger than the amplitude of scintillation fluctuations. So this effect deserves further study, but is unlikely to change results dramatically.

ORCID iDs

Daniel C. H. Gomes  <https://orcid.org/0000-0001-6299-2445>

Gary M. Bernstein  <https://orcid.org/0000-0002-8613-8259>

⁷ Extracted from the DECam exposure time calculator: <https://noirlab.edu/science/documents/scidoc0493>.

⁸ S. Gwyn (2024), private communication.

References

Andrae, R., Fouesneau, M., Creevey, O., et al. 2018, *A&A*, **616**, A8

Batygin, K., & Brown, M. E. 2016, *AJ*, **151**, 22

Bernardinelli, P. H., Bernstein, G. M., Abbott, T. M. C., et al. 2025, arXiv:2501.01551

Brown, M. E., & Batygin, K. 2016, *ApJL*, **824**, L23

Buie, M. W., & Keller, J. M. 2016, *AJ*, **151**, 73

Buie, M. W., Porter, S. B., Tamblyn, P., et al. 2020, *AJ*, **159**, 130

Elliot, J. L., Dunham, E., & Mink, D. 1977, *Natur*, **267**, 328

Euler, L. 1749, Mémoires de l'académie des sciences de Berlin, 3, 178

Fabry, C. 1929, *JO*, **12**, 1

Fernández, Y. R., Jewitt, D., & Ziffer, J. E. 2009, *AJ*, **138**, 240

Folkner, W. M., Williams, J. G., Boggs, D. H., Park, R. S., & Kuchynka, P. 2014, *IPNPR*, **196**, 42

Gladman, B., Lawler, S. M., Petit, J.-M., et al. 2012, *AJ*, **144**, 23

Gomes, D. C. H., Murray, Z., Gomes, R. C. H., Holman, M. J., & Bernstein, G. M. 2023, *PSJ*, **4**, 66

Harris, A. W., & Harris, A. W. 1997, *Icar*, **126**, 450

Hubbard, W. B., Hunten, D. M., Dieters, S. W., Hill, K. M., & Watson, R. D. 1988, *Natur*, **336**, 452

Hung, D., Tholen, D. J., Farnocchia, D., & Spoto, F. 2023, *PSJ*, **4**, 215

Kavelaars, J. J., Jones, R. L., Gladman, B. J., et al. 2009, *AJ*, **137**, 4917

LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201

Michell, J. 1767, *RSPT*, **57**, 234

Murray, Z. 2023, *PSJ*, **4**, 90

Neilsen, E. J., Annis, J. T., Diehl, H. T., et al. 2019, TM-2714-AE-CD-PPD, Fermi Lab

Nihei, T. C., Lehner, M. J., Bianco, F. B., et al. 2007, *AJ*, **134**, 1596

Ortiz, J. L., Sicardy, B., Camargo, J. I. B., Santos-Sanz, P., & Braga-Ribas, F. 2019, arXiv:1905.04335

Pannekoek, A. 1903, *AN*, **164**, 5

Pedersen, M. G., Escorza, A., Ppics, P. I., & Aerts, C. 2020, *MNRAS*, **495**, 2738

Petit, J.-M., Kavelaars, J. J., Gladman, B. J., et al. 2011, *AJ*, **142**, 131

Porter, S. B., Spencer, J. R., Verbiscer, A., et al. 2022, *PSJ*, **3**, 23

Pravec, P., Scheirich, P., Kušnírák, P., et al. 2006, *Icar*, **181**, 63

Rein, H., & Tamayo, D. 2016, *MNRAS*, **459**, 2275

Rice, M., & Laughlin, G. 2019, *AJ*, **158**, 19

Sheppard, S. S., & Trujillo, C. 2016, *AJ*, **152**, 221

Sicardy, B., Braga-Ribas, F., Buie, M. W., Ortiz, J. L., & Roques, F. 2024, *A&ARv*, **32**, 6

Sommargren, G. E., & Weaver, H. J. 1990, *ApOpt*, **29**, 4646

Tamayo, D., Rein, H., Shi, P., & Hernandez, D. M. 2019, *MNRAS*, **491**, 2885

Tran, T. X., Geller, S. R., Lehmann, B. V., & Kaiser, D. I. 2024, *PhRvD*, **110**, 063533

Trujillo, C. A., & Sheppard, S. S. 2014, *Natur*, **507**, 471